

THESIS: An important part of scientific activity consists in ascribing subjective probabilities to scientific hypotheses and changing these probabilities in the light of new evidence in accordance with Bayes's theorem: $p(H,E) = \frac{p(E,H)p(H)}{p(E)}$.

OBJECTIONS AND REPLIES

- O1. Subjective probabilities must be explicated in terms of betting odds; no bet on the truth of a scientific hypothesis could ever be settled; therefore one cannot ascribe subjective probabilities to scientific hypotheses. (De Finetti)
- R1. I reject the first premise. I regard the betting odds explication as merely a convenient operational definition (applicable only in certain idealized special situations) of the theoretical term 'subjective probability'. (Cf. Mach's definition of 'mass'.)
- O1a. One needs the betting odds explication in order to justify the probability axioms.
- R1a. I consider the probability axioms to have an analytic status with respect to our intuitive concept of probability.
- O1b. One needs the betting odds explication in order to attribute a definite numerical value to the subjective probability associated with any specific hypothesis.
- R1b. No. Consider a sequence of evidential statements implied by the hypothesis and such that if they were true the hypothesis would then be considered 'as likely as not' - i.e. $p(H,E) \approx 1/2$; estimate $p(E)$ by the betting odds explication; $p(E,H) \approx 1$; deduce $p(H)$ from Bayes's theorem.
- O2. Neither introspection nor R1b yields exact numerical probabilities for hypotheses.
- R2. I do not demand more than approximate or qualitative subjective probability ascriptions to scientific hypotheses. Indeed, given the unlimited amount of potentially relevant evidence that could be taken into account, any realistic predicting machine could do no better given a finite computation time: how closely it delimited a given subjective probability assignment would depend on the computation time. I claim the same is true of us.
- O3. All scientific hypotheses are universal generalizations over infinite domains; every universal generalization over an infinite domain must be ascribed a probability of zero (Popper, Carnap)
- R3. I reject both these claims. The infinity of space and time and their infinite divisibility are all scientifically contentious and certainly not to be understood as implicit in every scientific hypothesis. The domain of most hypotheses (e.g. that of Copernicus) is vague but not infinite. Furthermore outside Carnap's own idiosyncratic systems there is no objection to assigning finite probabilities to universal generalizations over infinite domains: indeed this can even be done on the basis of the betting odds explication if we define $p((x)Fx) = \lim_{N \rightarrow \infty} p(F1.F2.F3....FN)$.

- O4. Scientists prefer the simplest hypothesis, i.e. the least probable. (Popper)
- R4. If we have to choose between a circle and a particular ellipse, both consistent with the data, the former will typically get the greater subjective probability (and they are equally refutable). If the choice is supposed to be between a circle and some more general hypothesis (such as that of a general conic section) we cannot say which scientists would prefer (or why they should have to make a choice) without more information.
- O5. The most probable hypothesis will always be the one which goes least beyond the evidence; but scientists will always prefer the hypothesis of greatest scope, i.e. the least probable. (Popper et al.)
- R5. A scientist only has to choose between hypotheses with respect to some domain in which they make conflicting predictions. If one of the hypotheses makes successful predictions in some other domain within which the others make no predictions, then, other things being equal, it will end up with a greater posterior probability than the others for its predictions within the original domain. It will therefore be preferred within that domain.
- O6. A scientist may consider it worthwhile investing a lot of time and effort on a particular hypothesis (or research programme) for reasons independent of its (subjective) probability.
- R6. Not denied.
- O7. If we are interested in choosing the most informative experiments to carry out, they are not those which merely test the most probable hypothesis (or those that merely test the least probable one).
- R7. Agreed. Indeed they are precisely those that maximize the expected information gain, and this can only be evaluated by assigning subjective probabilities to the hypotheses.
- O8. Scientists tentatively accept certain hypotheses while tentatively rejecting others but they do not ascribe partial degrees of belief (i.e. subjective probabilities) to them. (Post ?)
- R8. No. We could show this by eliciting the scientist's potential betting odds with respect to specific predictions. If he really operated merely on the basis of an acceptance/rejection strategy his betting behaviour would be so deviant that we could reasonably expect to make a lot of money off him.
- O9. In order to make the subjective probability approach to scientific inference work, we have to be able to give an exhaustive list of all hypotheses which might eventually come to be preferred. But it is impossible to do this. (Hesse, Putnam, Shimony)
- R9. We can always do this provided we end up with a rag-bag of not-yet-thought-of hypotheses, to which we ascribe a small but finite prior probability. If this eventually gets a large posterior probability we can then dissect its contents further.
- O9a. We cannot psychologically go on assigning smaller and smaller subjective prior probabilities to more and more unlikely looking hypotheses indefinitely. We must therefore have a cut-off somewhere. But there will be trouble if later evidence ever comes to point strongly in favour of some hypothesis to which we initially ascribed a zero prior probability. (Hacking)
- R9a. No. The technique mentioned in R1b will enable us to ascribe finite probabilities (when we so choose) to any hypotheses however absurd that we ever could later come to favour ~~through~~ ^{through} a sequence of unlikely observations. Thus in general there will not be any lower bound to the subjective probabilities we are psychologically capable of assigning.

- O10. The probability axioms do not themselves imply that one must change probabilities in the light of evidence in accordance with Bayes's theorem. Nor does the Dutch Book argument, even if it were applicable to the case of scientific hypotheses. (Hacking, Hesse, et al.)
- R10. The first claim is doubtful: if we build a principle of total evidence into the interpretation of the axioms it is not clear that the claim is correct. The second claim is doubly false. Paul Teller has published a Dutch Book argument due to Putnam and David Lewis which does just this. I have extended it to cover the case of scientific hypotheses on which we cannot bet directly by considering bets on two successive predictions.
- O10a. Changing probability assignments in the light of evidence in accordance with Bayes's theorem is sometimes demonstrably irrational. (Mellor, Gillies)
- R10a. Rubbish.
- O11. There is no reason why assignments of probabilities to scientific hypotheses should satisfy the constraints imposed by the Dutch Book argument because we are not betting against Nature, Nature is not a vicious bookie trying to make money off us. (Putnam, Hesse)
- R11. This objection seems peculiarly muddled. It is not Nature but other blokes that could potentially take advantage of incoherent subjective probability assignments to hypotheses (provided we could use them to generate incoherent subjective probability assignments to sets of predictions). The objectors might equally well argue that the probabilities we assign to particular horses winning particular races need not be coherent because the horses are not trying to make money off us. Nature is the horse not the bookie.
- O12. If evidence favours a hypothesis and the hypothesis yields a particular prediction it is fallacious to infer that the evidence increases the probability of the prediction; $p(e_2, e_1)$ is a function of its arguments alone and is independent of the probabilities we ascribe to hypotheses: therefore we must add something else to subjective probability assignments and Bayes's theorem to do justice to actual scientific inference. (Hesse)
- R12. The first statement is true but not paradoxical (consider two hypotheses both favoured by the evidence in question but disagreeing about the prediction). The next two statements are false. There are no problems provided we always consider the effect of the evidence on all the serious rival hypotheses.
- O13. The subjective probability approach to scientific inference is only illuminating if we can prove a suitable theorem to the effect that however much scientists' prior probability assignments differ their posterior assignments will tend to converge with increasing evidence. But Savage's convergence-of-opinion theorem is far too restrictive to cover the typical scientific case. (Hesse)
- R13. I have proved a theorem which does I think meet the case.
- O14. Convergence of opinion depends on the rival theorists agreeing about the $p(e, h)$'s and merely differing about the $p(h)$'s. But how so when both are supposed to be equally subjective ($p(e, h) =_{\text{def}} p(e, h)/p(h)$)?
- R14. It seems to me a mistake to take this as a definition. If h is statistical $p(e, h)$ is not subjective but its value can be deduced from h . If we are

concerned with non-statistical hypotheses (as in most cases of scientific inference) then $p(e, h)$ would always be either 0 or 1, were h made sufficiently specific. But it is impossible to make h sufficiently specific (assuming e is not especially theory-laden) without listing an indefinitely large number of other conditions that must hold (subsidiary hypotheses about the behaviour of the apparatus etc., to take the most obvious example). $p(e, h)$ is therefore really to be thought of as a weighted sum of 0's and 1's where the weights are the $p(h')$'s for an unlistable number of subsidiary hypotheses h' .

If we look at things this way it's not surprising that two theorists could differ widely over their assignments of prior subjective probabilities to h and its immediate rivals, without there being equal disagreement over the $p(h')$'s. In fact if there were any obvious source of disagreement over the latter it could easily be ironed out by making the h' in question explicit and restoring agreement over the $p(e, h)$'s.

Q15. Bayes's theorem is only applicable when the evidence statement E becomes certain. This never happens in science. (R. Jeffrey)

R15. If there were significant doubt about E , we could always remove it by retreating to a more protocol protocol language. It seems simpler to adopt this strategy than to adopt Jeffrey's suggestion of a more complicated alternative to Bayes's theorem, allowing for a 'posterior' probability of E not equal to 1, since we couldn't be any more sure of such a probability ascription's value than we could otherwise have been of the truth of E .

Q16. The subjective probability approach to scientific inference may let in counter-inductive strategies. (Salmon)

R16. Specific counter-inductive inferences will be justified by virtually any theory (if I have sufficiently good reasons for believing that half the (finitely many) marbles in the urn are black, then the more white ones I have drawn out, the more probable that the next one drawn will be black). But I know no argument to suggest that the subjective probability approach will allow any general counter-inductive strategy.

Q17. Any theory of scientific inference worth its salt will solve the problem of induction (Popper), explain the success of science (Hesse).

R17. The subjective probability approach does lead us to expect our predictions to become increasingly successful. But of course it does not guarantee it in any particular case or over any particular period of time. One can imagine worlds in which the Bayesian would do badly in comparison with someone who always fastened on the least probable hypothesis as the basis for his next prediction. But the fact that we all ascribe negligible prior probabilities to such worlds is sufficient to make it rational for us to adopt a strategy which will work in all those worlds to which we ascribe non-negligible prior probabilities. It is always rational to act in accordance with one's best beliefs.

Q18. Scientists will often prefer a hypothesis that is actually inconsistent with some of the data. (Lakatos, Feyerabend, Kuhn)

R18. Taken in ^{that makes this true} ~~this~~ ^{subjective} sense, the 'data' don't get a probability 1 of proving correct.

In fact on the subjectivist account, provided a scientist's prior probabilities $p(h)$ are sufficiently deviant from those of other scientists, he can remain wholly rational in his inferences and yet behave in as crankish and Feyerabendian a manner as he likes. There is a lot of room for the subjective element in science. It is only scientific argument that is reduced to calculation.